IBM Fusion

Overview of an engineered Infrastructure

for

Red Hat OPENSHIFT

and some use cases

Beat Galbier
Storage Solution Architect – IBM Fusion
beat.galbier@ibm.com

IBM

Agenda

01 Why Red Hat OpenShift?

02 Infrastructure current state and challenges

03 IBM Fusion – the basics

04 IBM Fusion – Multi-Cluster

05 IBM Fusion – Control Tower

IBM

Why Red Hat OpenShift?

Red Hat named a “Leader” in the 2023 Forrester Wave™: Multicloud Container Platforms

“Red Hat sets the pace of the market with enterprise IT capabilities and massive market presence.”

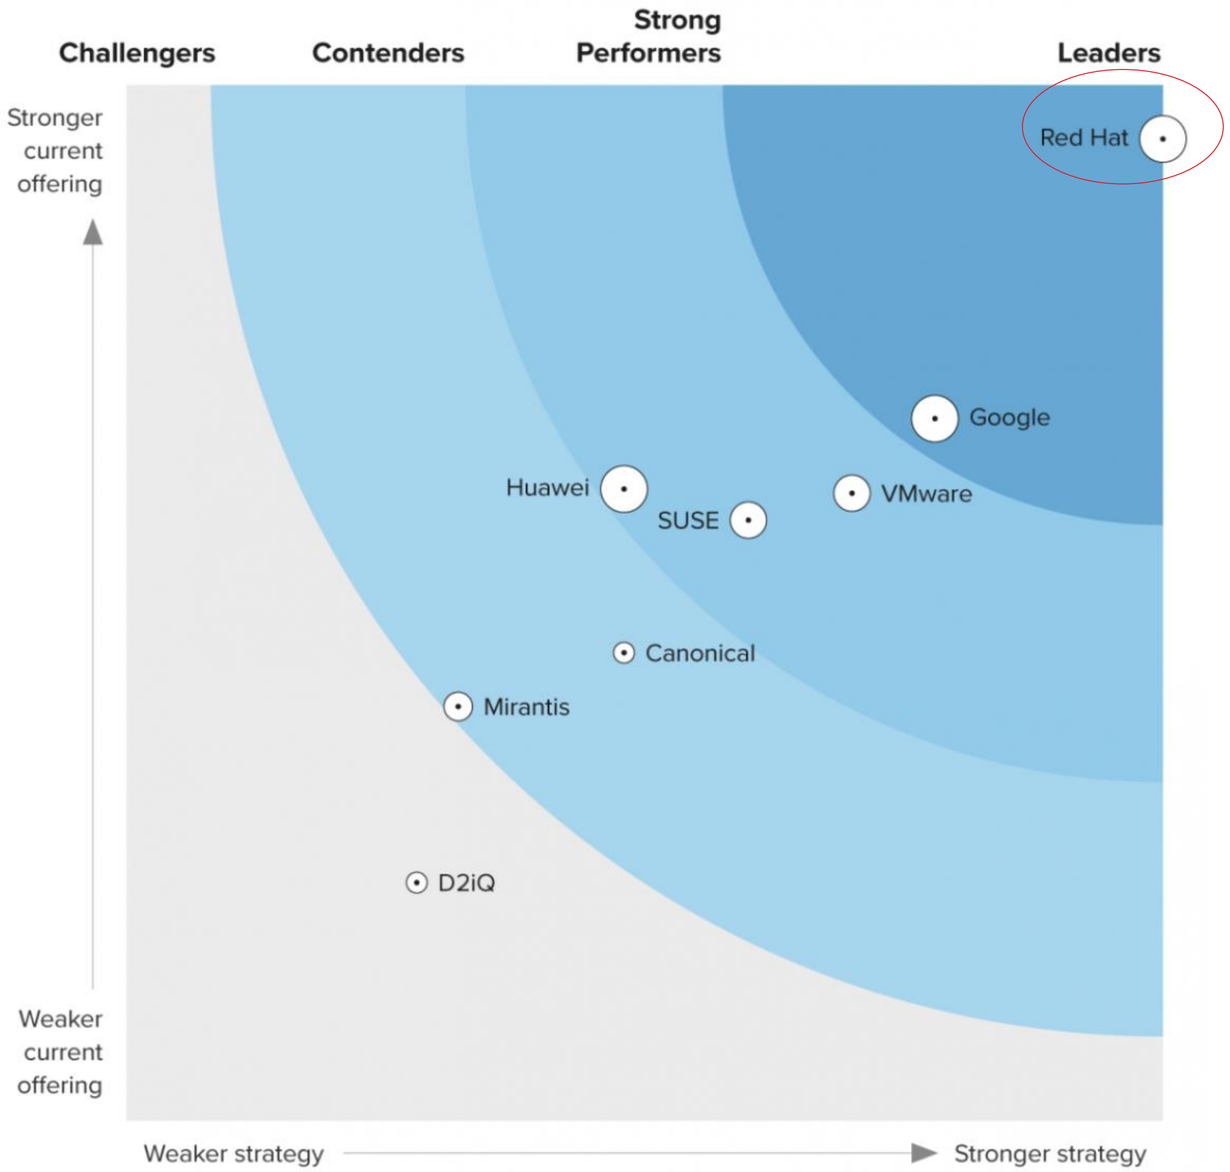

Consistent, everywhere!

Evolution and Challenges with current IT Infrastructure

vmware by Broadcom $$

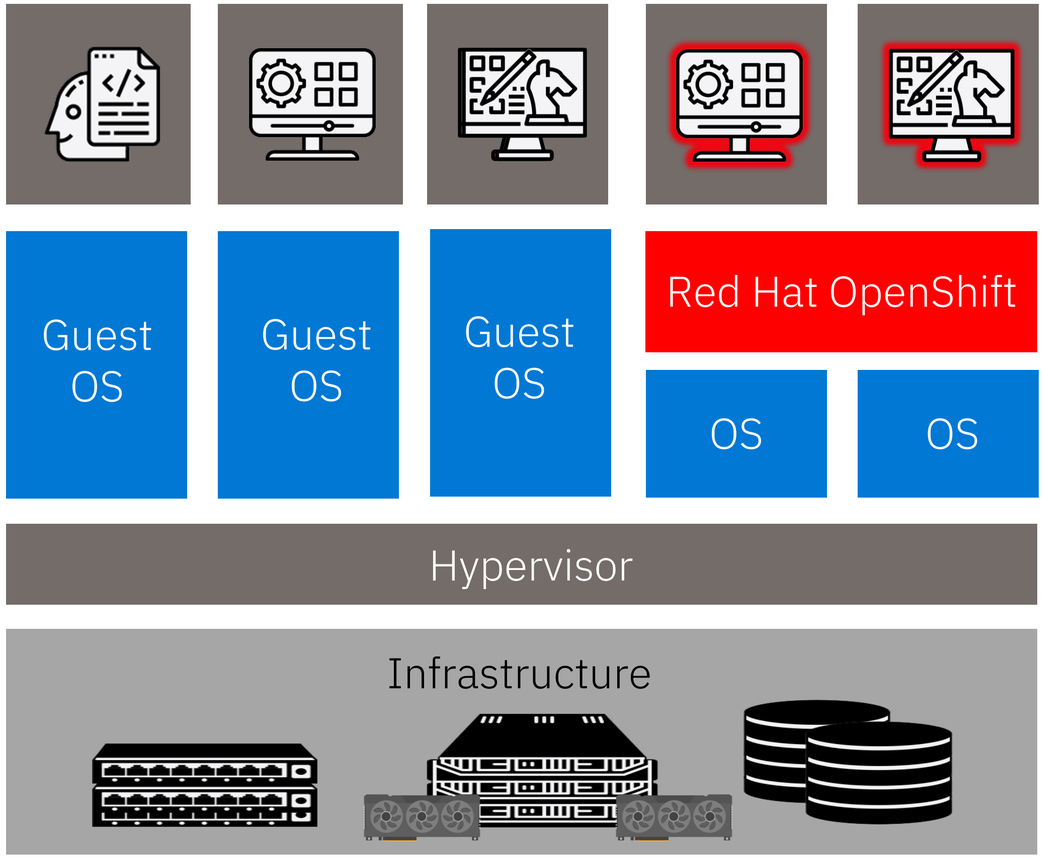

- Hypervisor overhead
- non-container-native storage
- network bandwidth
- scalability issues

$$

- additional Hypervisor licenses
- expensive Container Platform subscriptions
- 2 platforms to manage

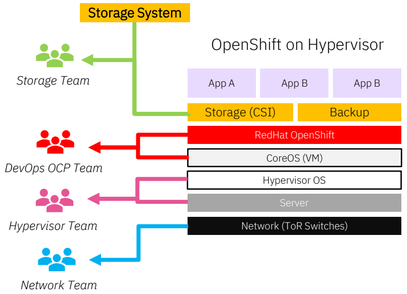

New Requirements

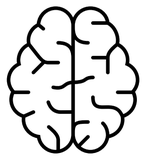

Solution Option A – Splitting Environments
go bare-metal OpenShift

vmware by Broadcom $$

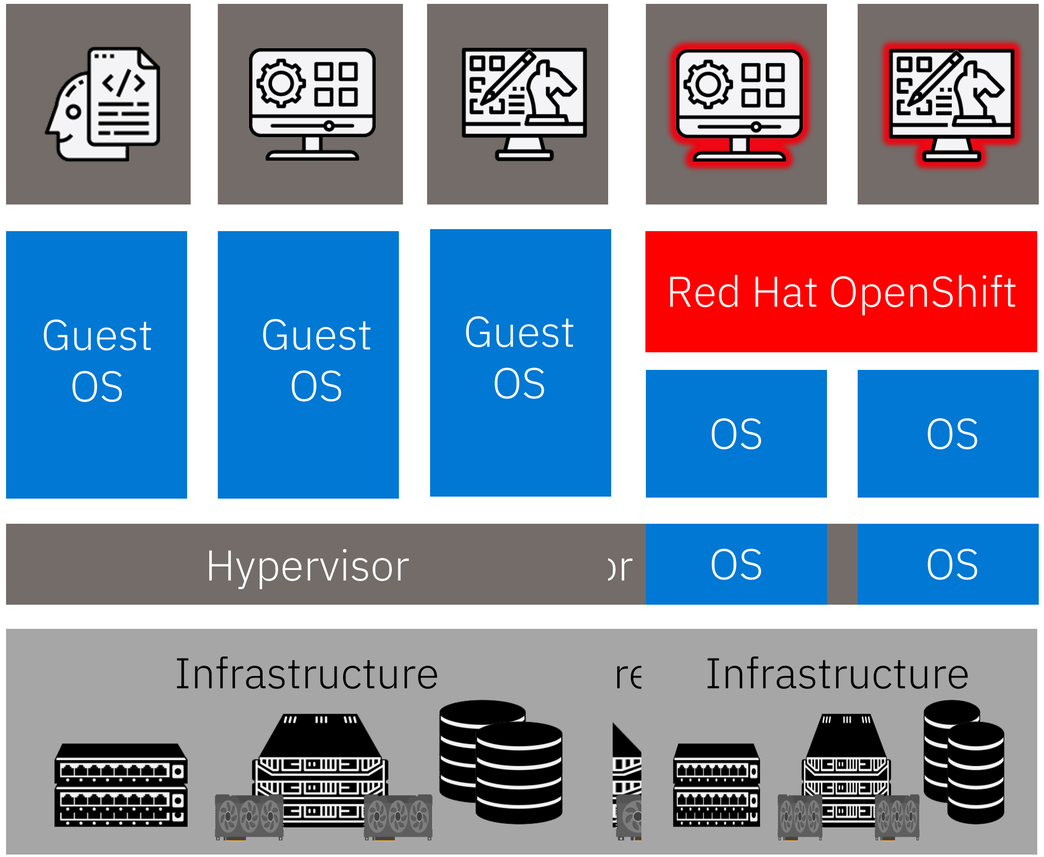

Siloed Infrastructures

- Hypervisor overhead ✓
- non-container-native storage ?
- network bandwidth ?
- scalability issues ?

$$

- additional Hypervisor licenses ✓
- Expensive Container Platform licenses ✓
- 2 platforms to manage

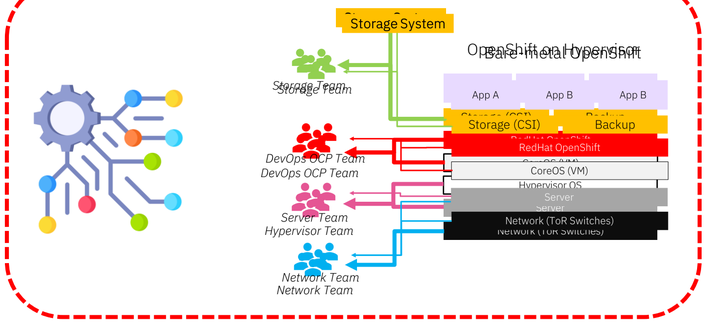

Bare-metal OpenShift isn’t easy…

OpenShift running on physical servers

- This software stack has multiple layers with multiple UIs / CLIs
- Multiple vendors are involved for support and licensing
- The software packages and install/update instructions must be obtained from multiple sources
- Interoperability of the components must be ensured and tested (e.g., Backup Tool -> Storage CSI -> OCP version)
- In many organizations, multiple teams are involved. Timing of change management must be coordinated to avoid outages

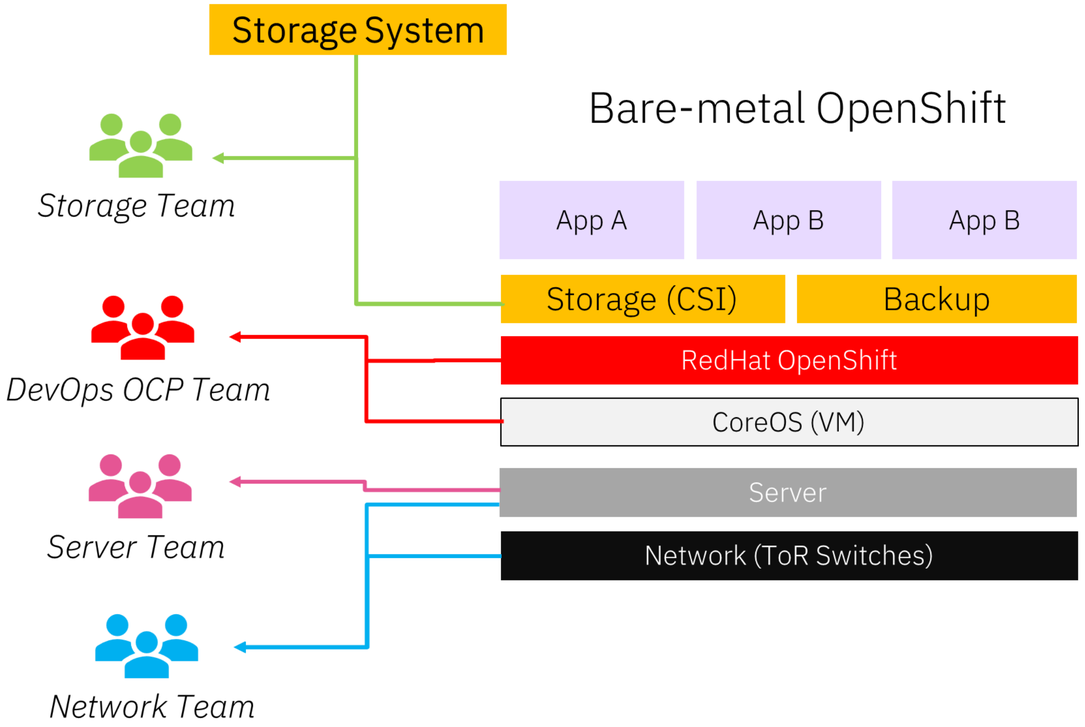

Bare-metal OpenShift – Installation Requirements

- Node requirements
- Planning a bare metal cluster for OpenShift Virtualization
- Firmware requirements for installing with virtual media
- Network requirements
  - Ensuring required ports are open
  - Increase the network MTU
  - Configuring NICs
  - DNS requirements
  - Dynamic Host Configuration Protocol (DHCP) requirements
  - Reserving IP addresses for nodes with the DHCP server
  - Provisioner node requirements
  - Network Time Protocol (NTP)
  - Port access for the out-of-band management IP address
- Configuring nodes
- Out-of-band management
- Required data for installation
- Validation checklist for nodes

Bare-metal OpenShift – Installation

- Installing RHEL on the provisioner node
- Preparing the provisioner node for OpenShift Container Platform installation
- Checking NTP server synchronization
- Configuring networking
- Establishing communication between subnets
- Retrieving the OpenShift Container Platform installer
- Extracting the OpenShift Container Platform installer
- Optional: Creating an RHCOS images cache
- Setting the cluster node hostnames through DHCP
- Configuring the install-config.yaml file
    - Configuring the install-config.yaml file
    - Additional install-config parameters
    - BMC addressing
    - Root device hints
    - Setting proxy settings
    - Configuring host network interfaces
    - Configuring host network interfaces for subnets
    - Configuring host network interfaces for dual port NIC
    - Configuring multiple cluster nodes
    - Configuring managed Secure Boot
- Manifest configuration files
    - Creating the OpenShift Container Platform manifests
    - Configuring NTP for disconnected clusters
    - Configuring network components to run on the control plane
    - Deploying routers on worker nodes
    - Configuring the BIOS
    - Configuring the RAID
    - Configuring storage on nodes
- Optional: Creating a disconnected registry
    - Preparing the registry node to host the mirrored registry
    - Mirroring the OpenShift Container Platform image repository for a disconnected registry
    - Modify the install-config.yaml file to use the disconnected registry
- Validation checklist for installation
- Deploying the cluster via the OpenShift Container Platform installer
- Following the installation
- Verifying static IP address configuration

+ Persistent storage integration

Introducing IBM Fusion HCI – an engineered Application Platform

or OpenShift in a Box

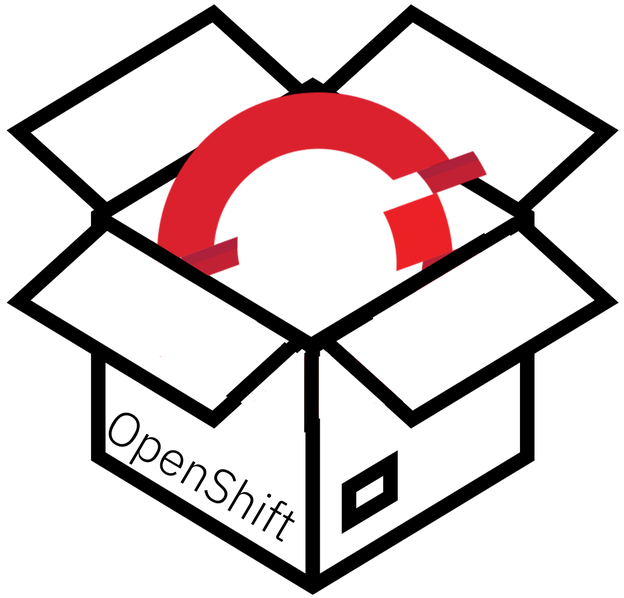

IBM Fusion is the modern platform to simplify providing and consuming modern cloud and VM applications consistently and dynamically

If it runs on OpenShift, it runs on Fusion | All IBM Cloud Paks | Databases | Off the shelf | Custom apps | watsonx

Fusion Software

Simple. Runs anywhere and everywhere

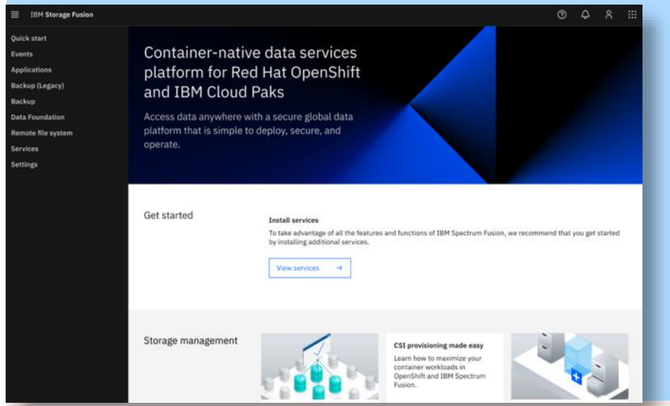

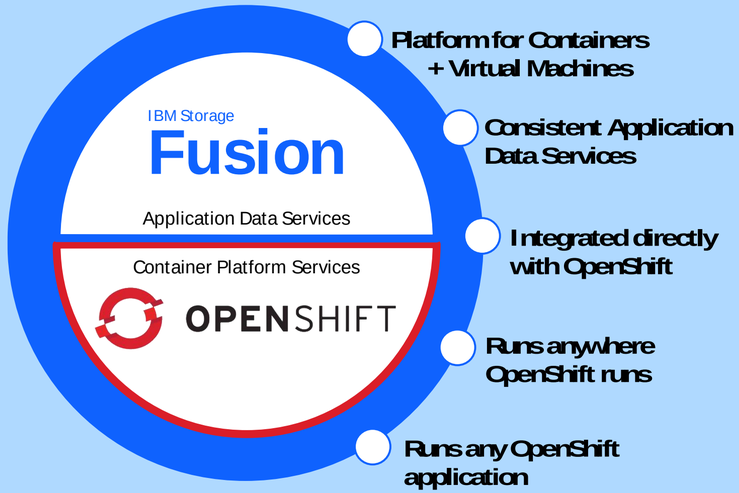

Simplified support

OPENSHIFT

aws | Microsoft Azure | IBM Cloud | vmware | IBM Z | IBM Power | Bare metal

Fusion HCI

OpenShift appliance

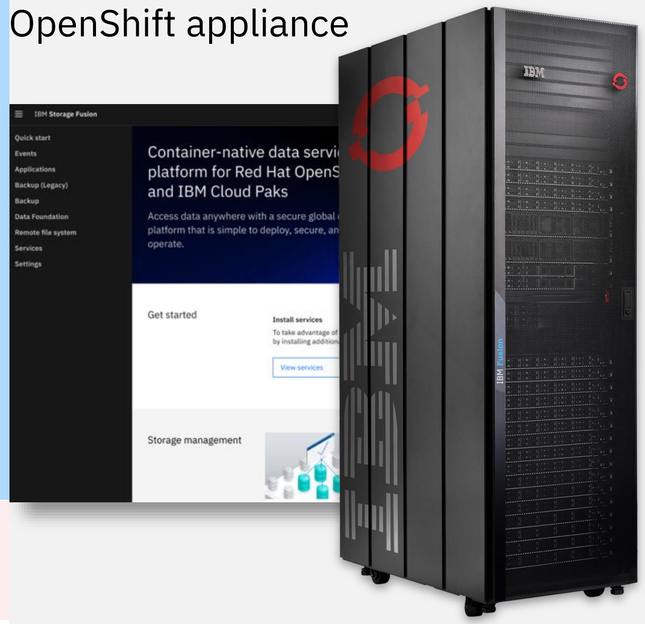

Engineered Infrastructure for OpenShift

IBM Fusion container-native data services

Portable to all clouds, operationally consistent everywhere

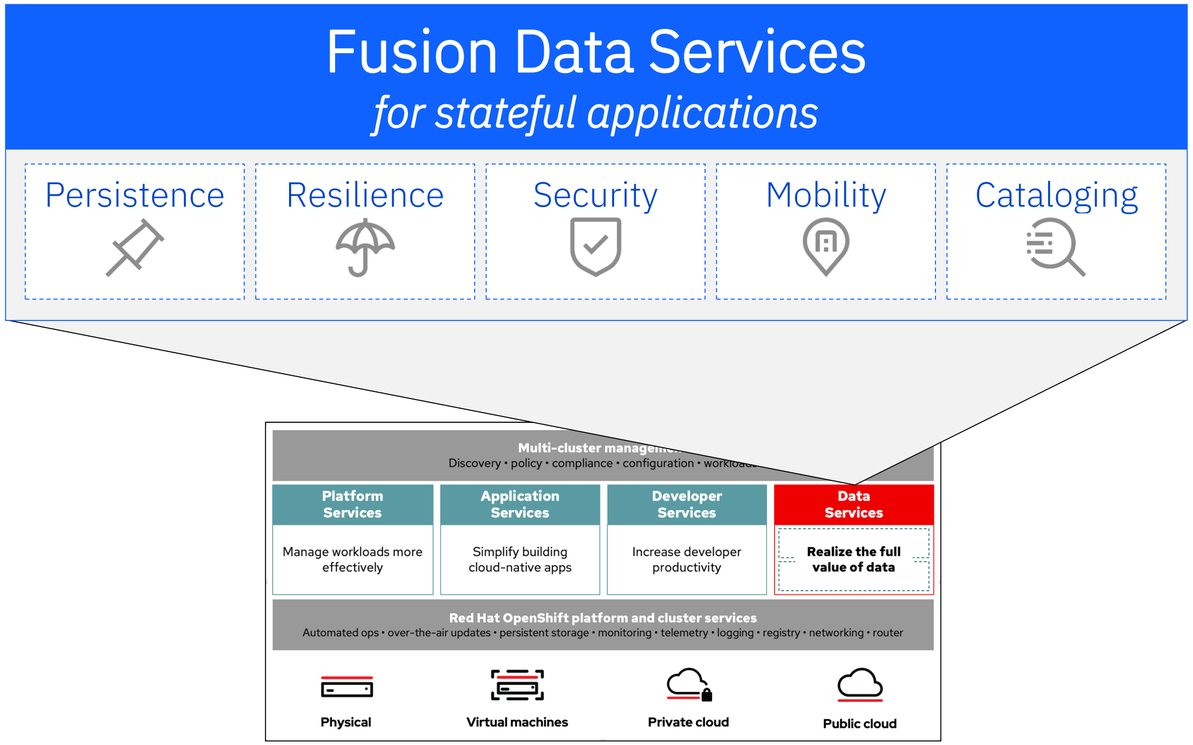

- **Persistent Volumes**
  CSI compliant file, block, and object
- **Backup and Restore**
  From simple snapshots to complex application consistent orchestrated backups. Ensure compliance with automated policy management.
- **Application HA/DR data**
  Synchronous or asynchronous data replication with tunable RPO and RTO
- **Remote storage acceleration**
  Transparently cache remote S3 data in the local file system to accelerate transactions
- **Enterprise-wide data catalog**
  Scan, index, and tag data residing anywhere in the organization for use in data science projects

IBM Fusion HCI - an engineered application platform

IBM Fusion HCI

Fusion data services

- Data Foundation
- Global Data Platform
- Backup & Restore
- Data Cataloging
- Disaster Recovery

Fusion base

- Serviceability
- GUI
- Observability

Infrastructure lifecycle management

- Compute management
- Network management
- Storage management

OpenShift Container Platform

Infrastructure

- Compute
- Network
- Storage
- GPU

IBM Fusion – Software

Complete Container Native Software Architecture

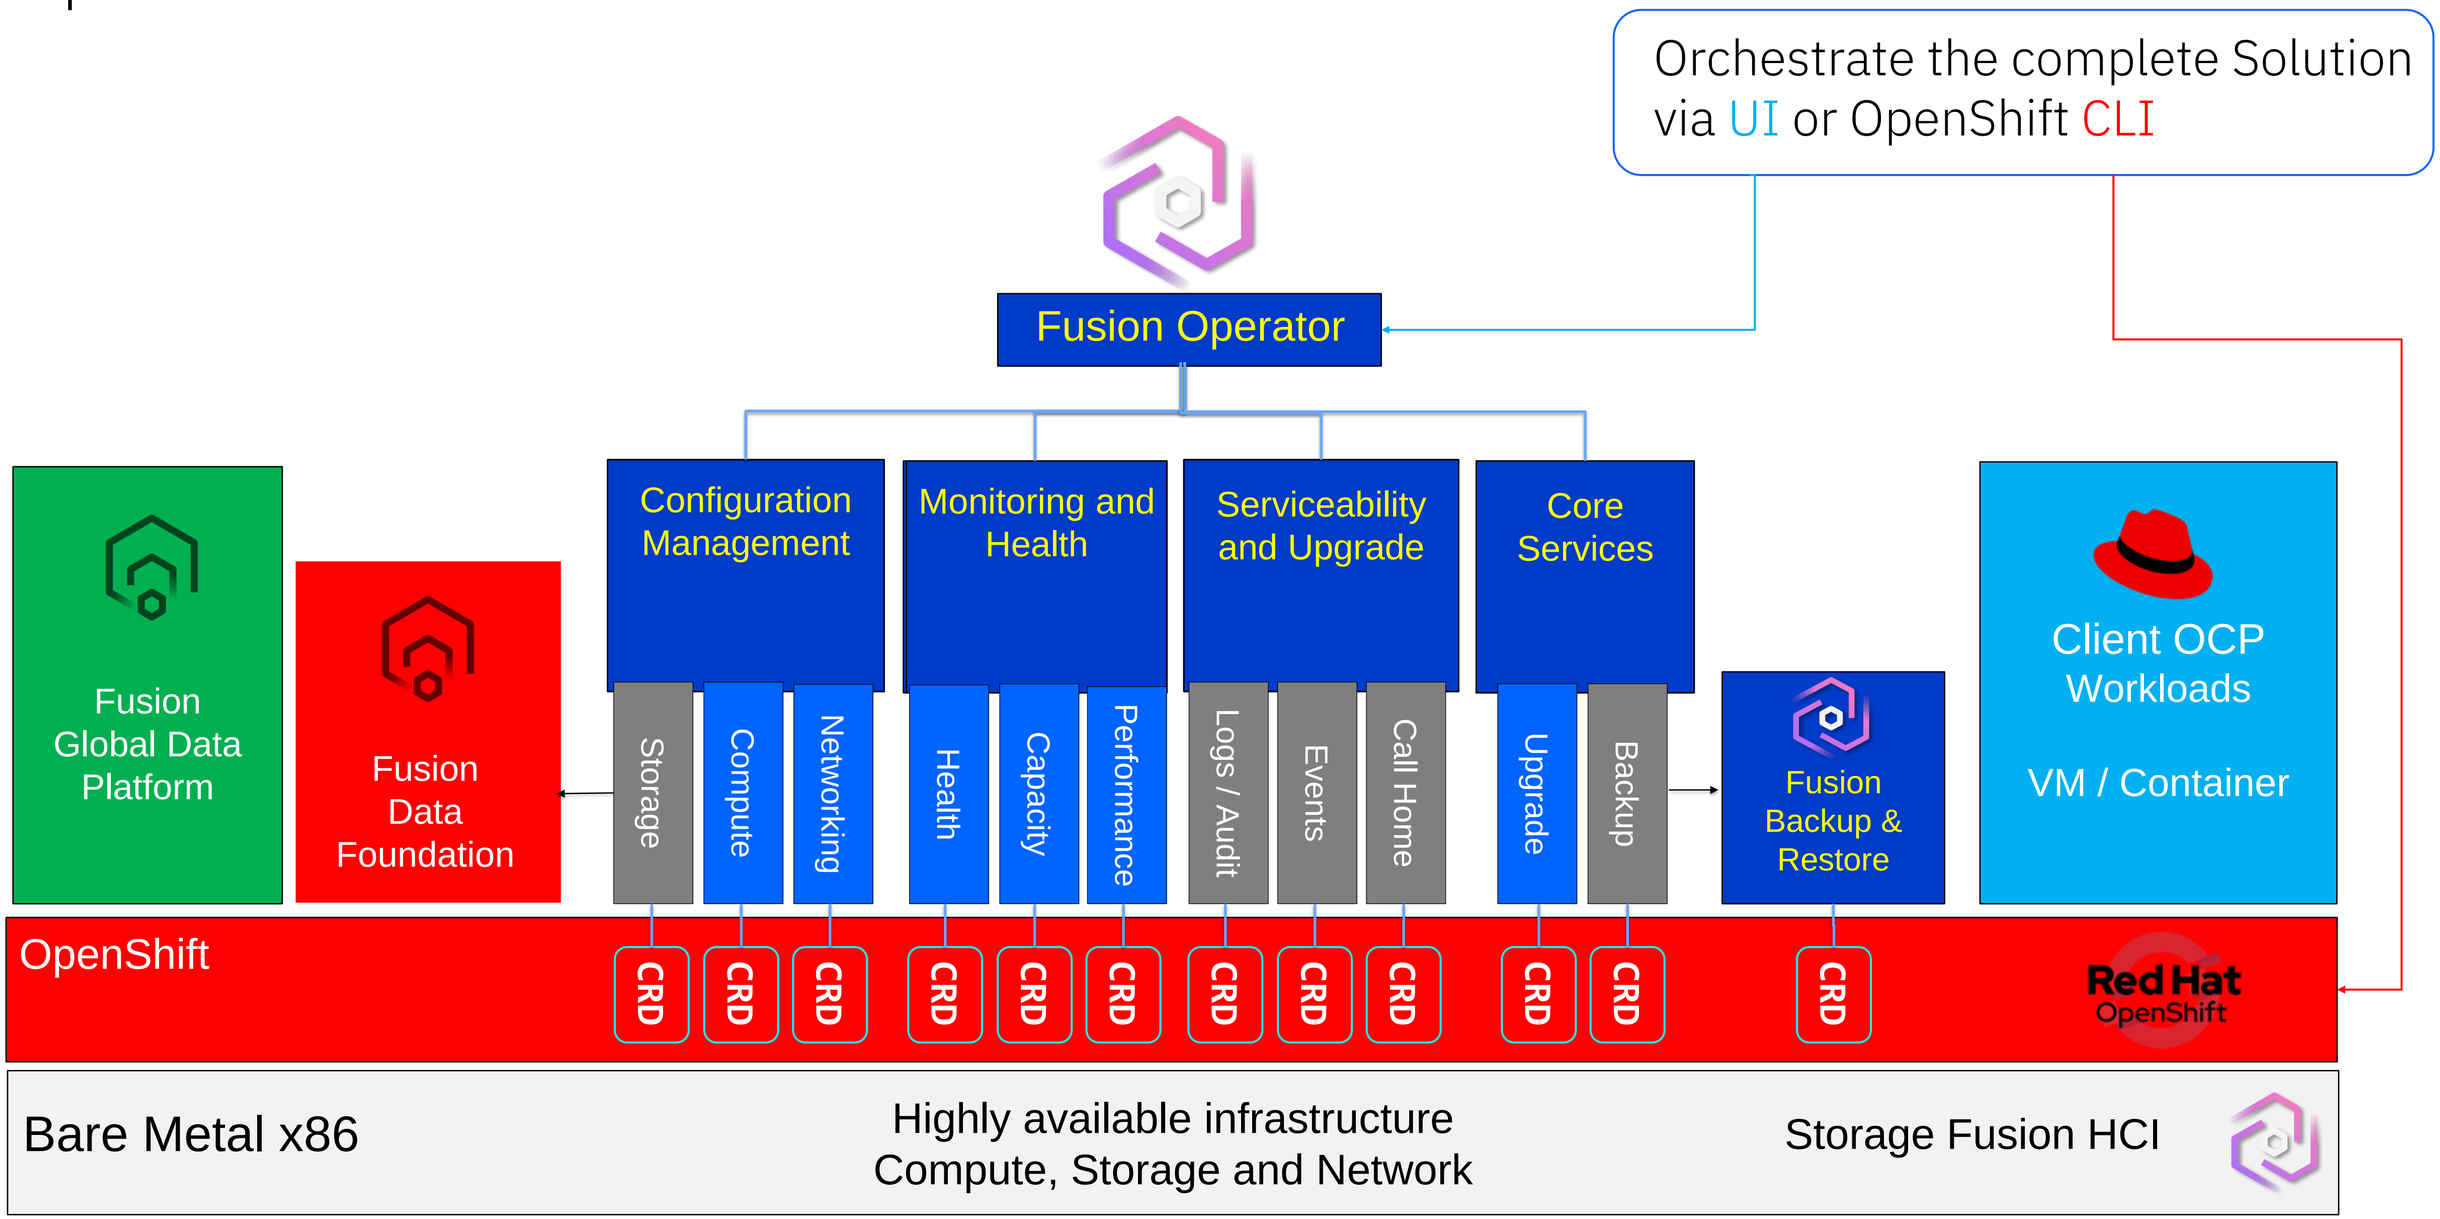

IBM Fusion HCI – Hardware

GPU nodes

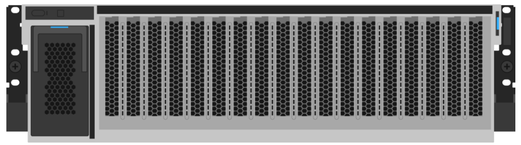

100 GbE OpenShift network

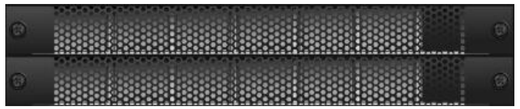

Management network

Hyper-converged servers with NVMe

GPU nodes
Up to 8x NVIDIA L40S or H100 NVL GPUs

High speed switches
100 GbE OpenShift network

High speed switches
Used for appliance management

Compute/storage nodes
Customizable worker nodes
Control nodes

AI

watsonx

Accelerate Generative AI Solutions
Inferencing, summarization, fine tuning

Data Lakehouse modernization

OPENSHIFT

Deploy in days vs months

App Modernization at your own pace
IBM Cloud Paks, ISV, or custom apps

Built-in HA/DR, Backup/Recovery

OPENSHIFT

Virtualization

VM Migration
Removing vSphere/vSAN dependency with container native storage

Built-in HA/DR, Backup/Recovery

IBM Fusion HCI – Scalability

6-node Fusion HCI 1 rack (min size)
- Raw storage with 12x 8 TB drives: 83 TiB
- Raw storage with 12x 4 TB drives: 41 TiB
- Total worker cores: 96 cores (192 vCPUs)

48-node Fusion HCI 3 rack (max size)
- Raw storage with 480x 8 TB drives: 3348 TiB
- Raw storage with 480x 4 TB drives: 1674 TiB
- Total worker cores: 2’880 cores (5’760 vCPUs)

16-node Fusion HCI 1 rack (max size)
- Raw storage with 160x 8 TB drives: 1116 TiB
- Raw storage with 160x 4 TB drives: 558 TiB
- Total worker cores: 832 cores (1’664 vCPUs)

Node configuration options
- 32 core or 64 core nodes with and w/o storage
- 256, 512, 1024, 2048 GB memory
- 4 TB or 8 TB NVMe drives
- up to 12 GPU nodes each with 1 to 8 GPUs – 96 GPUs

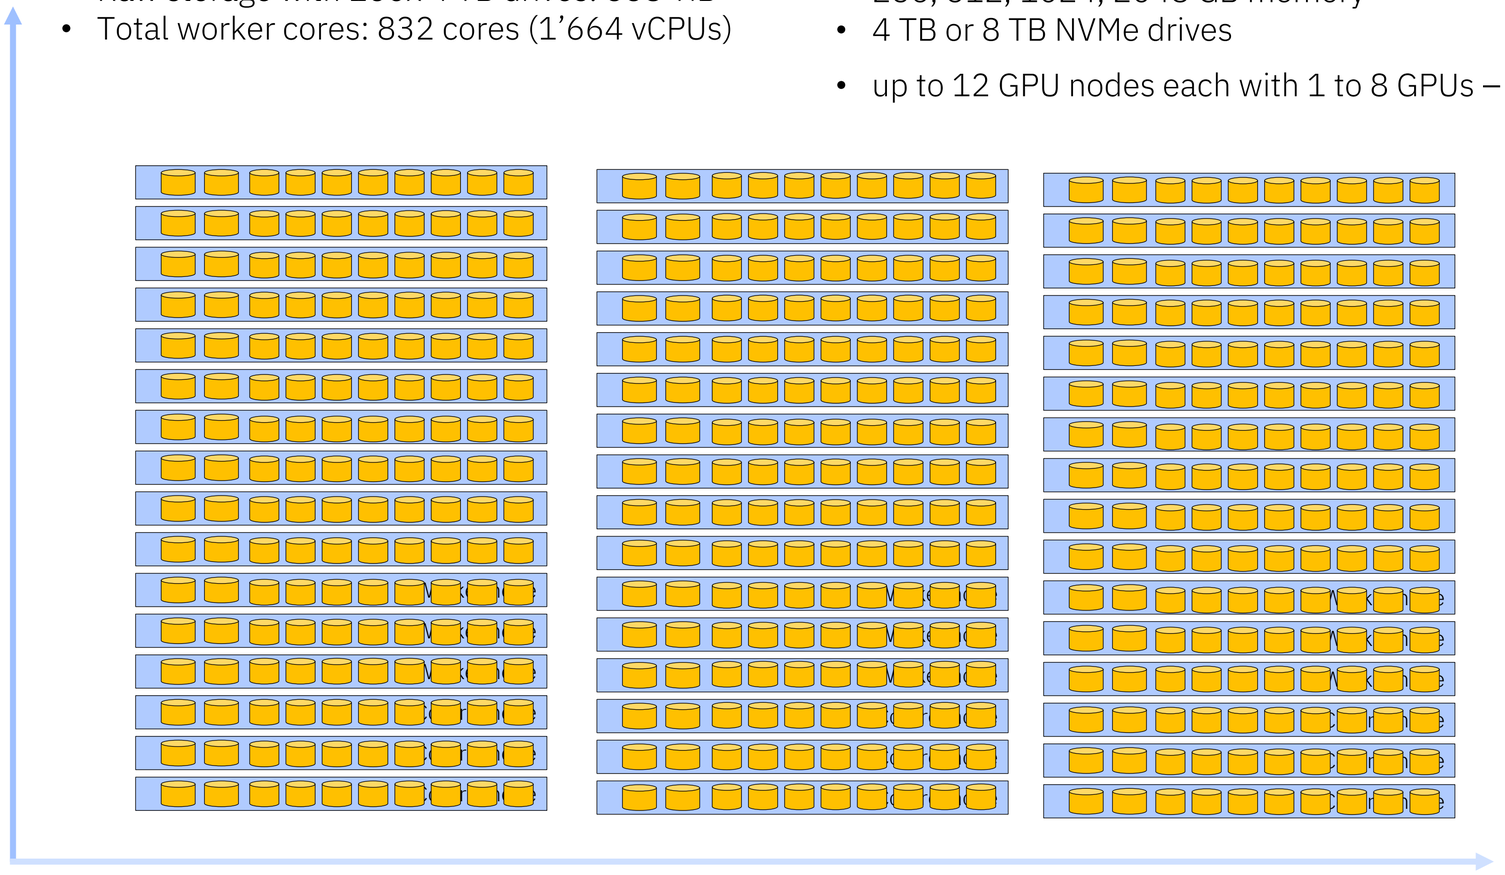

High availability and disaster recovery

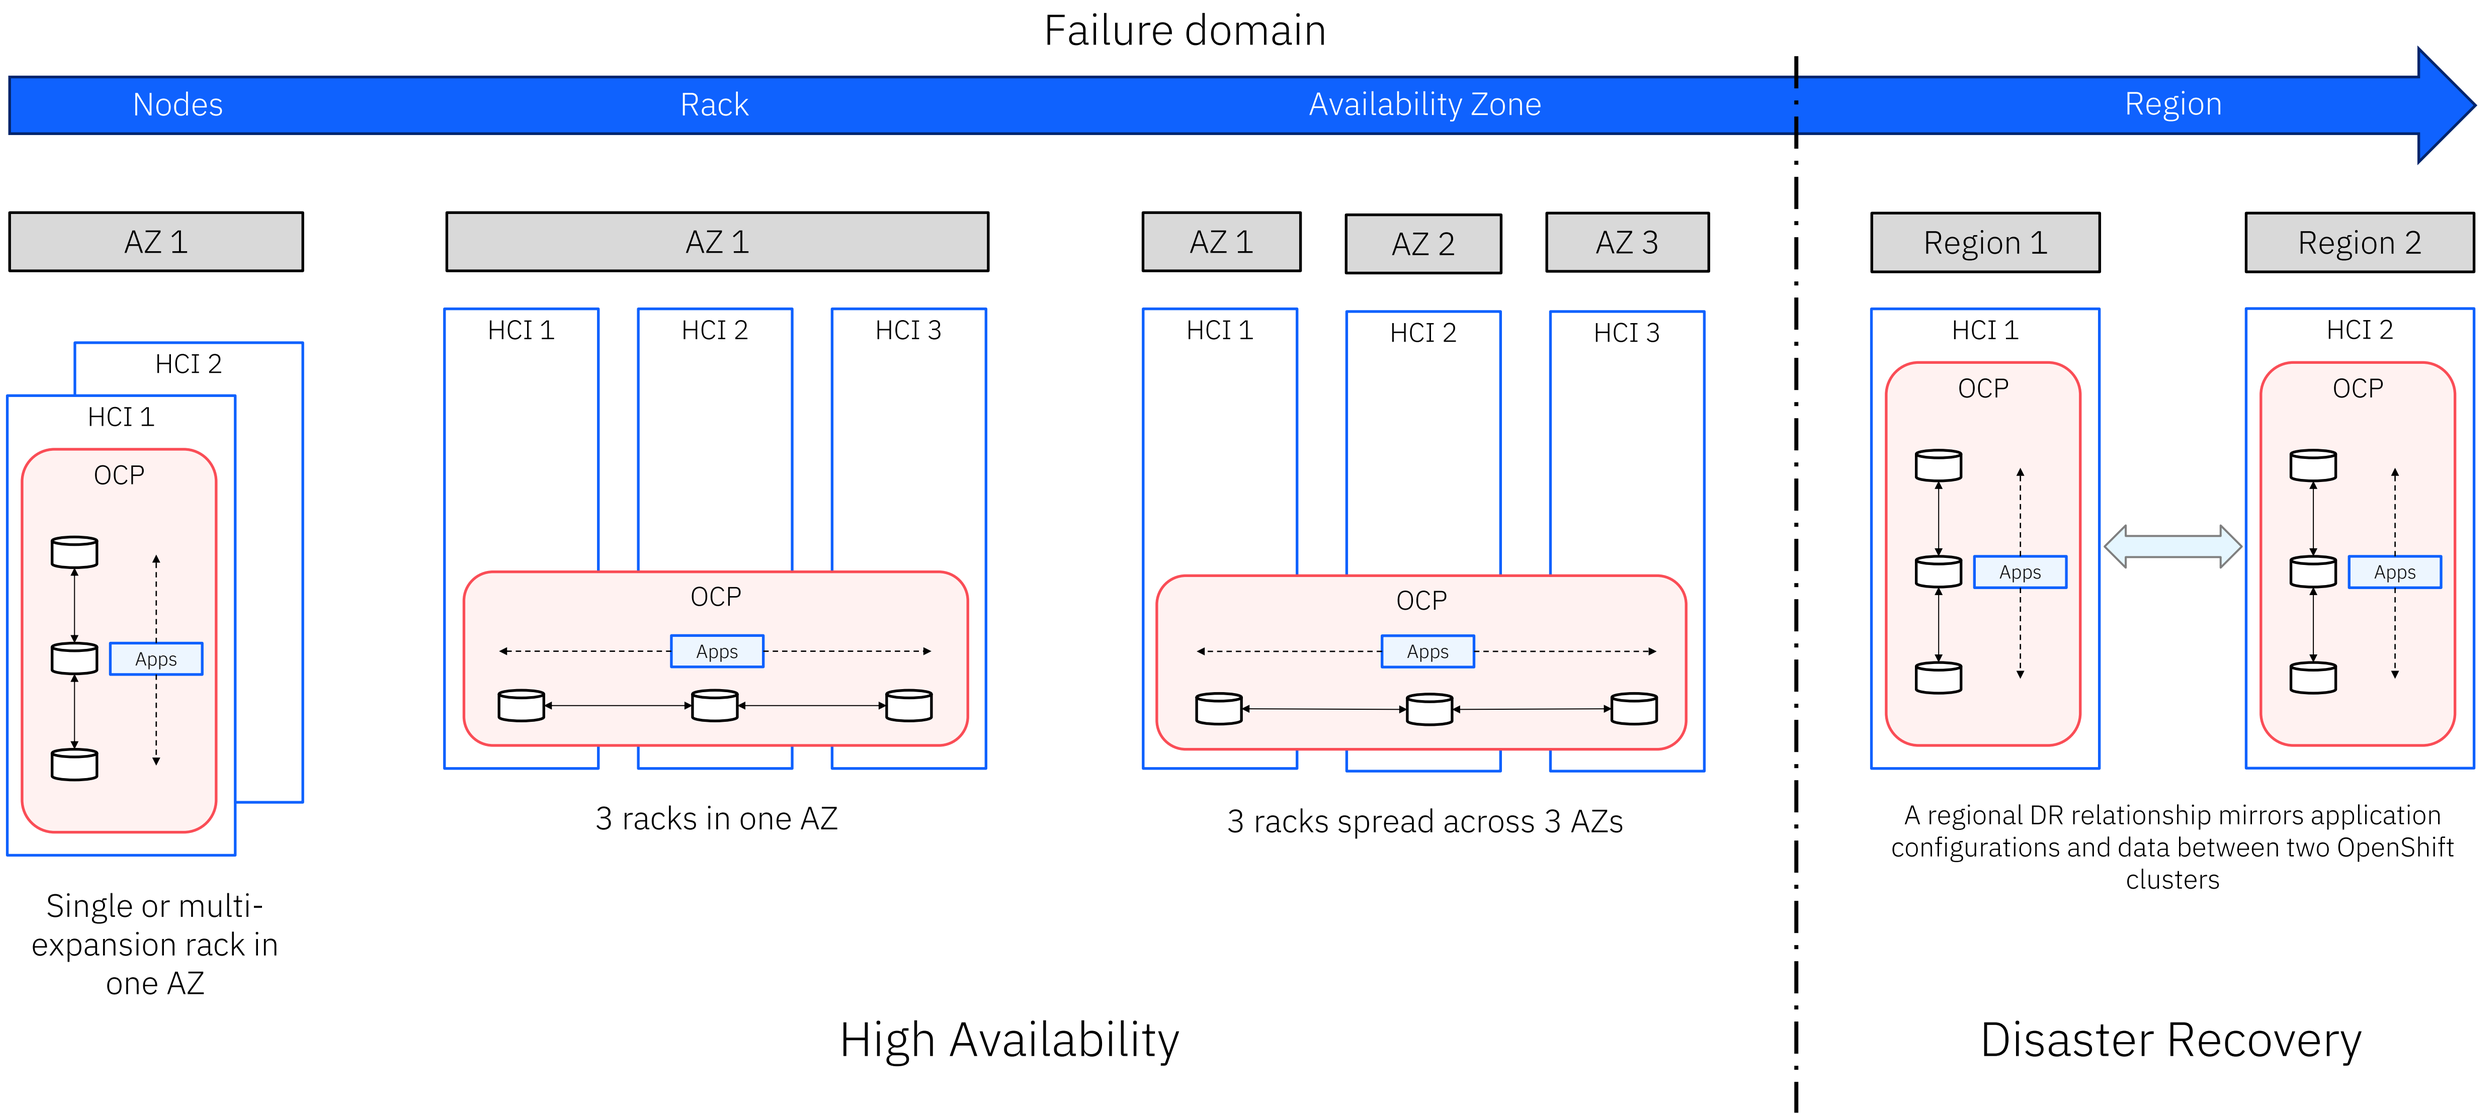

Remember? bare-metal OpenShift isn’t easy

OpenShift bare-metal DIY

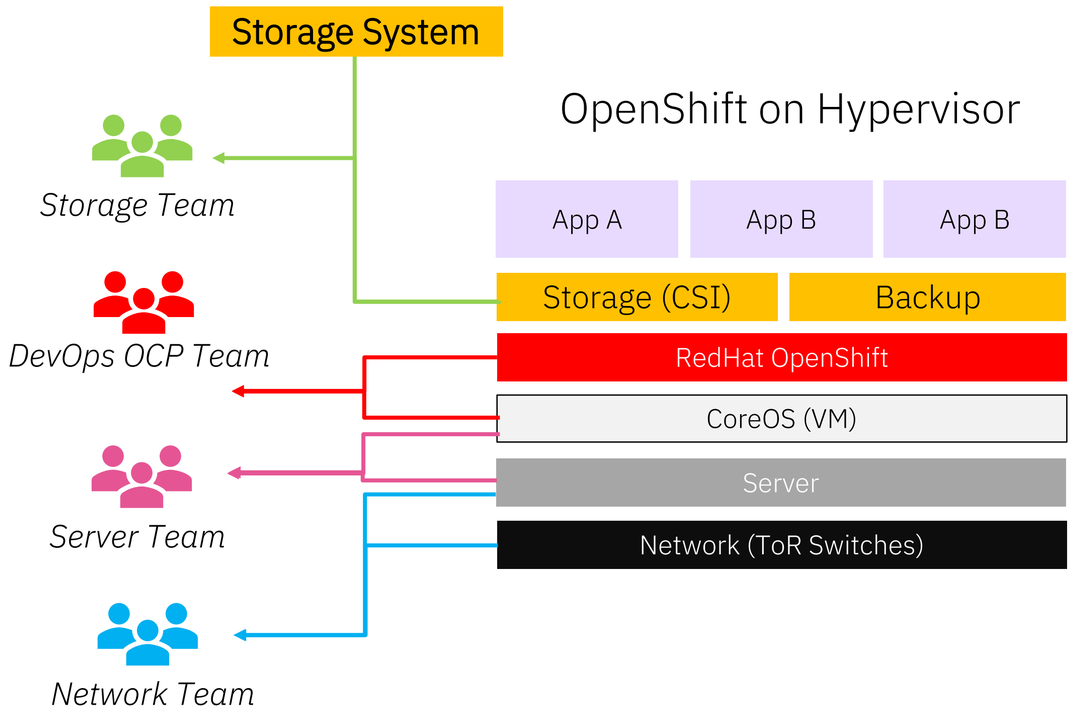

OpenShift bare-metal on IBM Fusion HCI

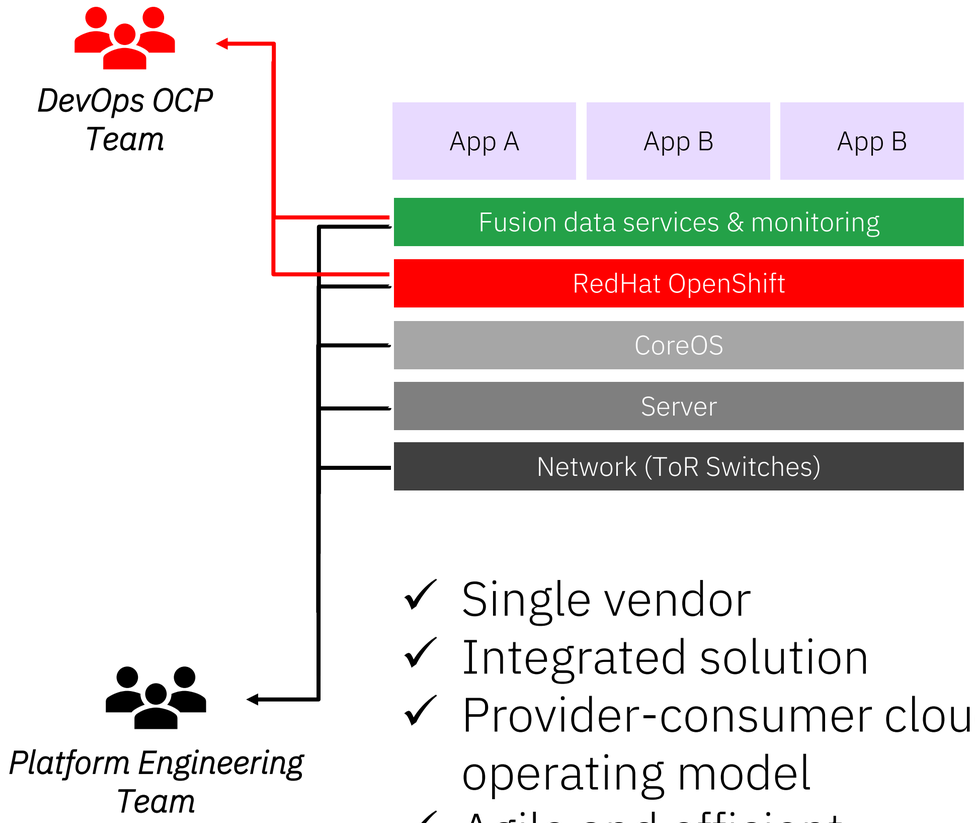

- ✓ Single vendor
- ✓ Integrated solution
- ✓ Provider-consumer cloud operating model
- ✓ Agile and efficient

Bare-metal OpenShift Cluster – Installation Requirements

with IBM Fusion HCI

- Network requirements
  - Ensuring required ports are open

  - DNS requirements
  - Dynamic Host Configuration Protocol (DHCP) requirements
  - Reserving IP addresses for nodes with the DHCP server

  - Network Time Protocol (NTP)

IBM will support the network planning. DNS, DHCP etc. configuration will be performed by the customer

Bare-metal OpenShift Cluster – Installation

with IBM Fusion HCI

- Optional: Creating a disconnected registry*
  - Preparing the registry node to host the mirrored registry
  - Mirroring the OpenShift Container Platform image repository for a disconnected registry

IBM will install/setup Fusion HCI and hand-over a running OpenShift cluster to the customer – ready for production

* Setup of registry is well documented and will be performed by the customer

Fusion HCI is the easiest infrastructure platform to deploy and manage OpenShift applications, VMs, and watsonx

Co-developed and tested for watsonx

Capabilities built for performance

- Efficient Cloud-operating model
  Highly automated, self-service and API centric
- Optimize cluster management
  (dev, test, experimental, production)
- Simplify maintenance Day 2
  life-cycle-management and simpler upgrades, scale out/up, monitoring, and health
- Complete HA/DR + application consistent backup services
- Linux & Windows Virtual Machines
  manage as containers in single control plane

Client Benefits

- Reduce total cost
  of OpenShift fleet management
- Developer and operational productivity
- Improve security
  managing all applications
- Reduce risk
  of missed schedules and SLOs
- Accelerate self-service cloud
  operating model on premises

Mission: Make it easy for platform engineers to design, deploy, and maintain infrastructure and services. Reduce deployment time from weeks to days.

IBM Fusion for **watsonx**

*Fusion is the easiest way to deploy and manage a self-service on-premises cloud for **watsonx***

Accelerate deployment and operations

- Reduce Day 1 deployment from months to days
- Accelerate and simplify operations with automated Day 2 monitoring, maintenance, and resiliency services
- Engineered for watsonx, integrated OpenShift system

Scale-out and scale-up easily

- Add GPU servers for watsonx.ai inferencing, prompt tuning, training
- Start data lakehouse on Fusion – grow with Ceph ready nodes
- Grow HCI with CPUs and Storage

Accelerate Performance

- Highest performing watsonx – better than Databricks at 60% cost
- Combined with watsonx, Fusion accelerates queries 7x to 90x (proven by IBM CIO office use case)

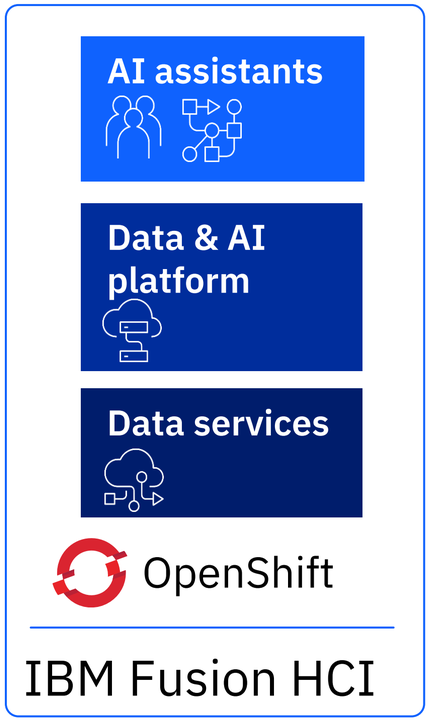

IBM Fusion HCI – the Evolution

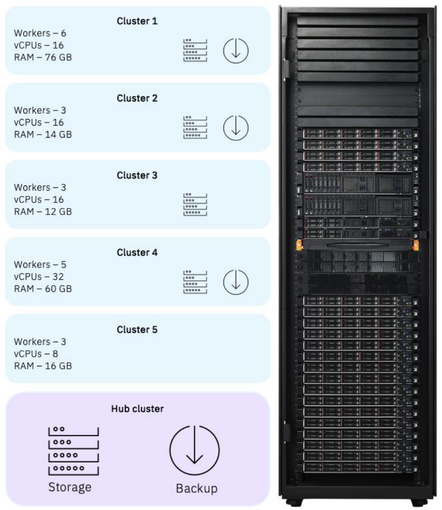

Single cluster HCI

Single OpenShift cluster

Quick time to value

Simplified Day 2 management

Resiliency – HA, B&R, DR

Multi-cluster HCI

Carve an appliance into multiple bare-metal and/or virtualized OpenShift clusters

Different OpenShift versions

Provision clusters on demand

Lower cost per cluster

OpenShift aaS on premise:

Based on bare-metal OpenShift with Hosted Control Planes

Managed hosted clusters provisioned on Fusion HCI System

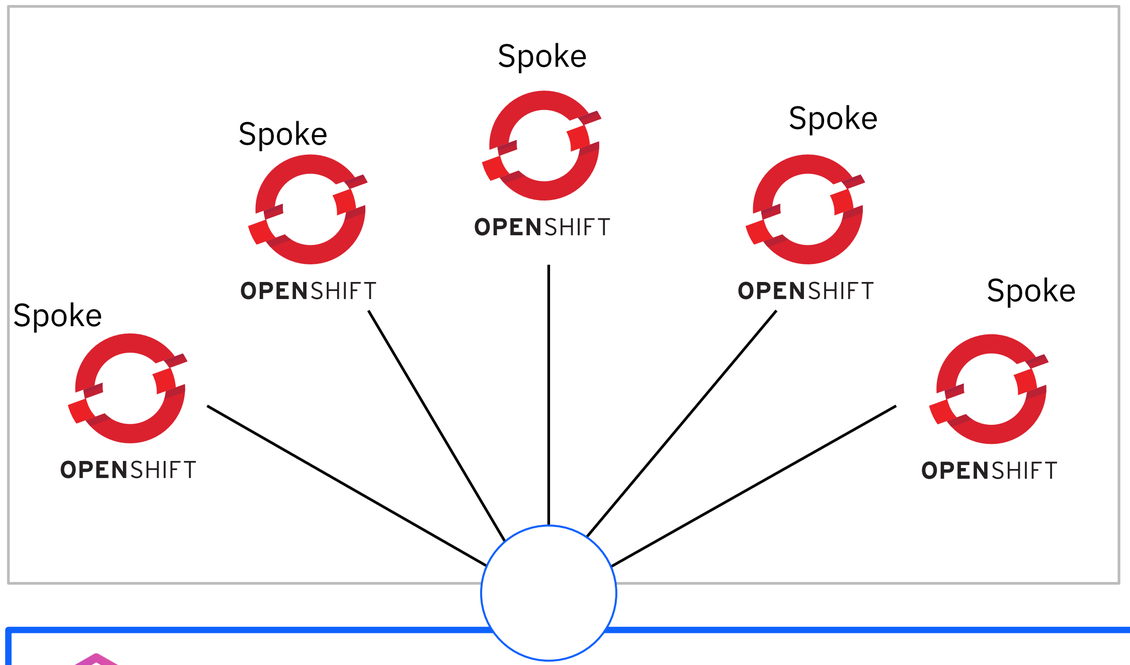

Application Developers

Site Reliability Engineers

Application DevOps Team

Interactions through APIs and Service Portals

Platform Engineering

Control planes

Backup engine

ACM

Storage

Fusion HCI System
Bare-metal Hub Cluster

Infrastructure	Services	Management
CoreOS	Storage (Block/File/Object)	Cluster Management
Servers	Backup & Restore	
Network switches	HA/DR	

IBM Fusion HCI – the Evolution

Single cluster HCI

Single OpenShift cluster

Quick time to value

Simplified Day 2 management

Resiliency – HA, B&R, DR

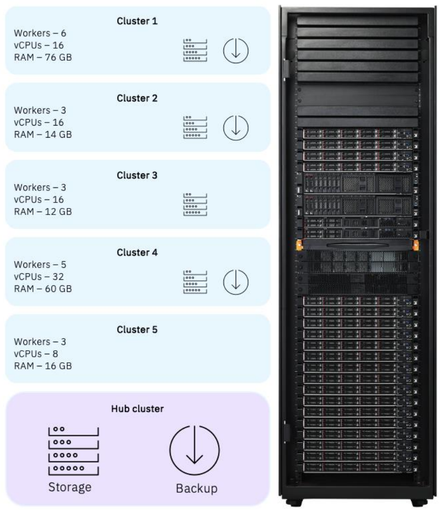

Multi-cluster HCI

Carve an appliance into multiple clusters

Different OpenShift versions

Provision clusters on demand

Lower cost per cluster

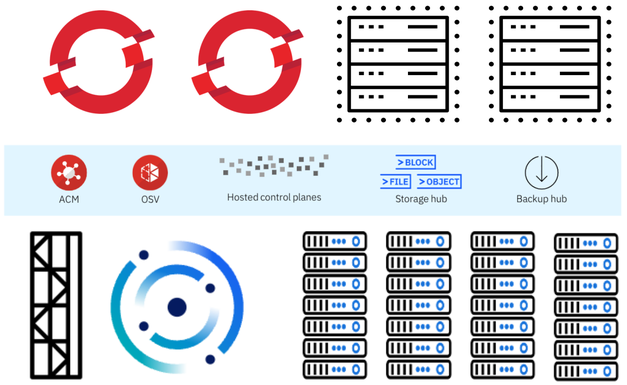

Hybrid cloud control plane for OpenShift

Extend the platform outside of the appliance

Take advantage of existing customer-managed bare metal inventory

Shift from running everything in the appliance to running on other infrastructure

IBM Fusion HCI – Control Tower
the complete picture

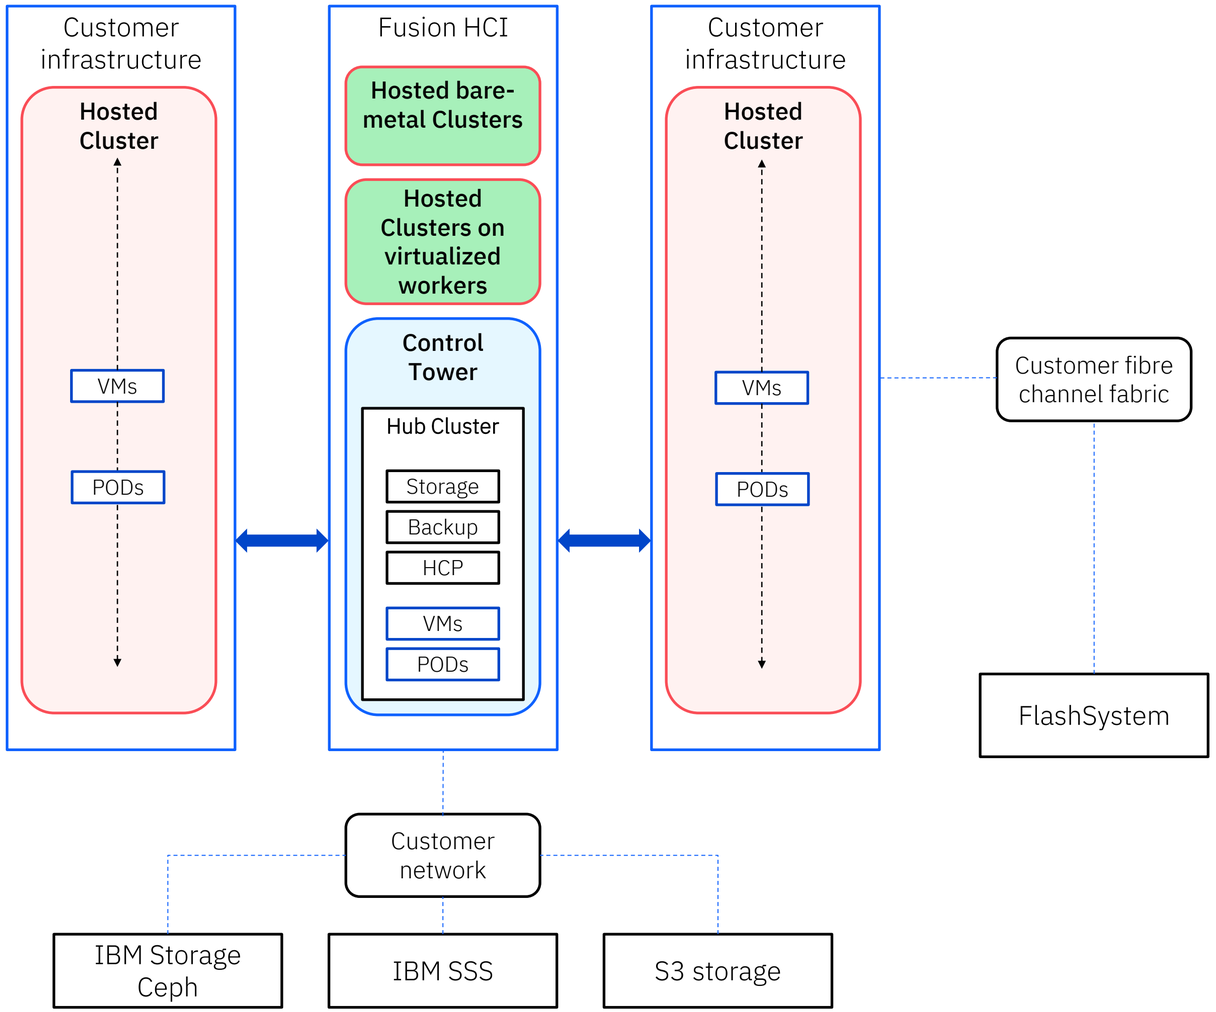

What benefit does Control Tower play?

- Separate the management of the OpenShift control plane from management of the hosted environment
- Storage hub provides storage classes (block, file, object) for the hosted cluster
- PODs and VMs in the hosted clusters are protected in the Backup Hub
- Manage multiple hosted clusters for different container and/or VM environments
- As legacy servers come off maintenance, they can be replaced with fully managed Fusion HCI servers

Thank you

But wait, there is one more thing I wanted to tell you...

Fusion HCI Workshop

For Customers

Join us at the IBM Innovation Studio in Munich for a Fusion HCI Workshop.

Highlights:

- Turnkey OpenShift experience
- Migrating VMs to OpenShift
- AI enabled Fusion HCI with GPUs
- OpenShift as a Service on premise
- Data Protection for OpenShift

Fusion HCI Workshop

For IBM Business Partners

Join us at the IBM Innovation Studio in Munich for a Fusion HCI Workshop.

Highlights:

- Turnkey OpenShift experience
- Migrating VMs to OpenShift
- AI enabled Fusion HCI with GPUs
- OpenShift as a Service on premise
- Data Protection for OpenShift
- Sizing
- Configuration (IBM Tools)
- TDA (Technical Delivery Assurance)

Fusion HCI Hackathon

For Enthusiasts

COMING SOON

Join us at the IBM Innovation Studio in Munich for a Fusion HCI Hackathon.

`OpenShift hands-on experience required!`

Hands-on Experience:

- Deploy your own OpenShift Cluster
- Install the Fusion Operators
- Deploy VMs on OpenShift
- Backup and Restore your Application

Available starting 3Q25